

ConvAbuse: Data, Analysis, and Benchmarks for Nuanced Abuse Detection in Conversational AI

Amanda Cercas Curry¹ and Gavin Abercrombie¹ and Verena Rieser^{1,2}

¹The Interaction Lab, Heriot-Watt University, Edinburgh

²Alana AI

{ac293, g.abercrombie, v.t.rieser}@hw.ac.uk

Abstract

We present the first English corpus study on abusive language towards three conversational AI systems gathered ‘in the wild’: an open-domain social bot, a rule-based chatbot, and a task-based system. To account for the complexity of the task, we take a more ‘nuanced’ approach where our ConvAI dataset reflects fine-grained notions of abuse, as well as views from multiple expert annotators. We find that the distribution of abuse is vastly different compared to other commonly used datasets, with more sexually tinted aggression towards the virtual persona of these systems. Finally, we report results from bench-marking existing models against this data. Unsurprisingly, we find that there is substantial room for improvement with F1 scores below 90%.

Warning: This paper contains examples of language that some people may find offensive or upsetting.

1 Introduction

Abusive language detection has received extensive attention for social media, (see e.g. Vidgen et al., 2020a), but far less within the context of conversational systems. As argued by UNESCO (West et al., 2019), detection and mitigation of abuse towards these (often anthropomorphised) AI systems is important in order to avoid reinforcement of negative gender stereotypes. Following this report, several recent works have investigated possible abuse mitigation strategies (Cercas Curry and Rieser, 2018, 2019; Chin and Yi, 2019; Ma et al., 2019). However, the results of these studies are non-conclusive as they are not performed with live systems nor with real users – mainly because of the absence of reliable abuse detection tools. The majority of currently deployed systems use simple keyword spotting techniques, (e.g. Ram et al., 2018; Khatri et al., 2018), which tend to produce a high number of false positives, such as cases in which the user expresses frustration, or use of profanities for

emphasis, as well as false negatives, e.g. missing out on subtler forms of abuse (Han and Tsvetkov, 2020). Recently, Dinan et al. (2019); Xu et al. (2020) released an abuse detection tool trained on Wikipedia comments and crowd-sourced adversarial user prompts (the latter are not freely available). Whereas in this work,

- We show that the distribution of abuse towards conversational systems is vastly different compared to other commonly used datasets, with more than half the instances containing sexism or sexual harassment.
- We develop and release a detailed annotation scheme with the help of experts.
- We use this scheme to annotate a corpus of 20k ratings on >6k samples (ca. 2k from each system), which we call ConvAbuse. We critically discuss and experiment with different labelling methods for this task. We also release a subset of 4k examples and their expert annotations.¹
- We benchmark commonly used abuse detection methods on this corpus.

2 Related work

Most work on detecting harmful content such as offensiveness, toxicity, abuse, and hate speech (see Fortuna et al. (2020) for definitions), has focused on social media platforms, foremost Twitter (e.g. Ball-Burack et al., 2021; Basile et al., 2019; Cao and Lee, 2020; Davidson et al., 2017; Fortuna et al., 2020; Founta et al., 2018; Gröndahl et al., 2018; Koufakou et al., 2020; Nejadgholi and Kiritchenko, 2020; Nozza et al., 2019; Razo and Kübler, 2020; Vidgen et al., 2020b; Wang et al., 2020; Waseem et al., 2017; Zampieri et al., 2019b,a, 2020), or other social media platforms,

¹We make our code and two of the datasets available at <https://github.com/amandacurry/convabuse>. We are unable to release the data from one system for privacy reasons.

including Facebook (Glavaš et al., 2020; Zampieri et al., 2020), Gab (Chandra et al., 2020), and Reddit (Han and Tsvetkov, 2020; Zampieri et al., 2020). Work has also been undertaken on data from comments on news media (Glavaš et al., 2020; Razo and Kübler, 2020; Wang et al., 2020; Zampieri et al., 2020), chatrooms and discussion forums (Gao et al., 2020; Wang et al., 2020), Wikipedia discussions (Fortuna et al., 2020; Glavaš et al., 2020; Gröndahl et al., 2018; Nejadgholi and Kiritchenko, 2020; Pavlopoulos et al., 2020), and message services such as WhatsApp (Saha et al., 2021).

Meanwhile, there has been relatively little work on abuse detection for conversational AI. Furthermore, much of the work that does exist in this area does not actually involve human-machine dialogue: Dinan et al. (2019); Xu et al. (2020) use a classifier developed on Wikipedia comments, which was further trained on adversarial prompts collected via crowd-sourcing. Similarly, de los Riscos and D’Haro (2021) designed a chatbot to intervene against online hate speech, trained and evaluated on data from Wikipedia and Civil Comments.

Those few studies that do report abuse detection results from genuine human-machine conversations tend not to include publicly released datasets. These include several submissions to the Amazon Alexa Challenge² (Cercas Curry et al., 2018; Khatri et al., 2018; Paranjape et al., 2020). As such, to the best of our knowledge, this is the first study to release a public dataset of human-machine conversations for the task in this domain.

While we aim to detect abuse directed against any target, gender-based abuse has been identified as a particularly prevalent problem in conversational AI (Cercas Curry and Rieser, 2018; Silvervarg et al., 2012; West et al., 2019), and abuse detection systems have themselves been found to contain gender biases (Park et al., 2018). Misogyny and sexism detection has been applied to social media in binary (Fersini et al., 2018; Nozza et al., 2019) and multi-class (Waseem and Hovy, 2016) settings. We extend this to take an intersectional approach, analysing multiple types of abuse in a hierarchical multi-label framework (see §3.3).

3 The ConvAbuse corpus

We collected data from conversations between users and three different conversational AI sys-

²<https://developer.amazon.com/alexaprize> (accessed May 2021.)

tems, which have different goals and properties. Two of them are classed as chatbots, i.e. social, open-domain systems, while the other is a transactional, goal-oriented system. The first two systems listed below are text-based, whereas the last system is voice-based with a synthetic female-sounding voice. As such, two out of the three systems are female gendered, either by voice or name.

Alana v2 An entrant to the Alexa Challenge 2018, a competition in which university teams develop social chatbots which aim to hold engaging conversations with users in the United States. The bot implemented a mixture of social chit-chat and provision of information via entity linking. Users were notified of the competition at the beginning of the conversation. We only have access to the automatically transcribed user utterances, which contain recognition noise. The data was collected between April 2017 and November 2018.

CarbonBot. An assistant created by Rasa³ and, hosted on Facebook Messenger.⁴ The bot aims to convince the user to buy carbon offsets for their flights. It also notifies the user that conversations will be recorded for research purposes. The data was collected between 1st October 2019 and 7th December 2020.

ELIZA. An implementation of the rule-based conversational agent intended to simulate a psychotherapist (Weizenbaum, 1966), designed for academic purposes, and hosted at the Jožef Stefan Institute.⁵ It aims to engage the users by asking open questions: “*Tell me more about <X>!*”. The data was collected between 19th December 2002 and 26th November 2007.

For example conversations from all three systems, see Appendix A.

3.1 Pre-processing

For each system, we discarded any test conversations involving the systems’ developers, and extracted the utterances from all user turns from the conversations. Following the findings of Pavlopoulos et al. (2020) that dialogue context can affect (and even reverse) human judgement of toxicity, we included the system output as well as the previous turns (where available) of both user and system.

³<https://rasa.com> (accessed May 2021.)

⁴<https://m.me/CarbonBot.from.Rasa> (accessed May 2021.)

⁵<http://www-ai.ijs.si/eliza> (accessed May 2021.)

We removed any system output that is not directly provided to the user in text form (such as voice prosody tags), and replaced web addresses with the token <URL>.

3.2 Sampling

Previous research has shown that 5-30% of user utterances are abusive (Cercas Curry and Rieser, 2018). In order to find these instances, one can use purposive nonprobability sampling using abusive keywords. However, this can lead to the creation of heavily biased datasets (Vidgen and Derczynski, 2021; Wiegand et al., 2019). We attempted to strike a balance between obtaining a high proportion of examples that contain abusive language and not biasing the datasets towards explicit forms of abuse that contain such keywords. To do this we combined two sets of keywords:

1. A list of ‘profanities’ — 265 regular expressions from a blacklist obtained from Amazon. These keywords are mostly profane, offensive words, which can be expected to capture use of explicitly offensive language.
2. 1,532 terms from Hatebase,⁶ — a crowd-sourced list of hate speech to capture (i) abuse targeted at specific groups such as women and racialised minorities, (ii) more subtle forms of abuse that do not contain explicitly offensive language, and (iii) terms that have taken on abusive meanings recently or in certain subcultures. As most of the terms also have other, non-hateful meanings (Sap et al., 2019), we hypothesised that their use as keywords could capture abusive content, while not biasing the data towards purely offensive terms.

We then used stratified sampling, to extract utterances at random from six stratas of the datasets that contained conversations featuring 0, 5, 10, 15, 20 and 25 per cent of sentences that feature terms from the list of keywords. As the total number of conversations and user turns in CarbonBot is smaller, we did not sample from this, annotating the entire dataset. We used the bias metrics of Ousidhoum et al. (2020), finding that the final corpus does not seem to be heavily biased towards typical abusive language keywords (for details, see Appendix B).

3.3 Annotation scheme and guidelines

We created a hierarchical labelling scheme based on insights from prior work. At the top level, we

adapted Poletto et al. (2019)’s unbalanced rating scale, in which input is labelled from +1 (*friendly*) to -3 (*strongly abusive*), providing information about not only whether or not it is considered to be abusive, but also the severity of any abuse:

- 3. Strongly negative with overt incitement to hatred, violence or discrimination, attitude oriented at attacking or demeaning the target.
- 2. Negative and insulting/abusive, aggressive attitude.
- 1. Negative and impolite, mildly offensive but still conversational.
0. Ambiguous, unclear.
1. Non-abusive.

Based on Waseem et al. (2017)’s two-dimensional typology of abuse, we then elicited labels for the *target* (*group*, *individual–system*, or *individual–3rd party*) and *directness* (*explicit* or *implicit*). To obtain more finely-grained information about the targets of abuse, annotators then label the instances as either *general*, *sexist*, *sexual harassment*, *homophobic*, *racist*, *transphobic*, *ableist*, or *intellectual*. These labels were based on known factors in the matrix of domination (Collins, 2002). These *type* classes are not mutually exclusive, allowing the annotations to capture intersectionality.

To allow for contextual interpretations, annotators were shown the target user utterance, the agent’s utterance to which it responded, and a previous speaking turn by both the user and the agent.

In supervised learning for text classification tasks, human-provided labels are typically aggregated to one ‘gold-standard’ label per instance by means of majority-vote, adjudication, or statistical methods. However, the notion of reducing multiple annotations to a single ‘correct’ label has been criticised for erasing minority perspectives (Blodgett, 2021; Gordon et al., 2021). This is because perception of phenomena such as hate, varies both across individuals and culturally (Salminen et al., 2018). We therefore retain and evaluate classification systems on the labels of all the annotators.

3.4 Annotators

We recruited eight gender studies students in their early 20s. Six of them identify as female, and two as non-binary. All are L1 English speakers, predominantly from the United Kingdom, except for one from the United States. One identifies as Asian, the remaining seven as white. Full details are provided in the data statement in Appendix D.

⁶<https://hatebase.org> accessed 6th Nov 2020.

3.5 Agreement measurement and analysis

We adjusted the annotation scheme iteratively in three rounds by observing the labels applied to batches of 100 random examples from the data. We measured agreement with Krippendorf’s α (α), which can take account for multiple annotators, missing values, and ordinal ratings (Gwet, 2014). Where agreement was low, we invited our experts to discuss examples. However, since abuse is a subjective phenomenon, we did not force agreement. We discarded the data used in guideline development, and the annotators labelled the rest of the data according to the final guidelines. Agreement scores per annotation task are shown in Table 1. Overall, the annotators achieved moderate to substantial agreement for the majority of categories. Agreement was consistent across datasets. We report on intra-annotator agreement in Appendix C.

Annotation task	Label type	Overall
Abusive/non-abusive	Binary	0.69
Abuse severity	Ordinal	*0.46
Type	Binary \times 8	0.79
Ableism	Binary	0.73
Homophobia	Binary	0.83
Intellectual	Binary	0.63
Racism	Binary	0.96
Sexism	Binary	0.63
Sex harassment	Binary	0.84
Transphobia†	Binary	0.00
General	Binary	0.74
Target	Nominal	0.61
Directness	Binary	0.26

Table 1: Inter-annotator agreement: Krippendorf’s α . *Ordinal weighted α used; † based on only 8 examples.

Sexism. Although the annotators form a fairly homogeneous group in terms of demographics and all have a background in Gender Studies, we find only moderate α for sexism, consistent with previous studies which found that up to 85% of disagreement was on this category (Waseem and Hovy, 2016). We find that sexism and sexual harassment are closely intertwined but distinct, with 47% of examples labelled sexist also judged to be sexual harassment but only around 22% of sexual harassment also being sexist. Some annotators see all sexual harassment as necessarily sexist as it is rooted in misogyny. This is in agreement with the European Centre for Gender Equality which states that ‘sexual harassment is an extreme form of sexism’.⁷ In

⁷<https://eige.europa.eu/publications/sexism-at-work-handbook/part-1-understanding/what-sexual-harassment> (accessed May 2021.)

our data, sexist examples focus on using gendered slurs such as “*bitch*”, and sexual harassment uses sex as a way to create a hostile and offensive environment though it may not contain explicit terms, e.g. “*I wanna see you naked*”.

Directness. Low inter- but moderate intra-annotator agreement (see Tables 1 and 10) suggests this task is highly subjective and open to interpretation. For example, annotators may perceive abuse as more implicit that is phrased as a question (e.g. “*are you stupid*”, “*can i be your lover?*”), that is misspelled/misheard, (e.g. “*Connie Lingu*”, or comments with sexual connotations but no overtly sexualised words (e.g. “*call me big daddy*”). Annotators can disagree not only on whether abuse is implicit or explicit, but whether it is abuse at all. Examples of disagreement between explicit abuse and non-abuse include commonly used expressions of frustration or surprise such as “*wtf*”. Implicit abuse is particularly difficult to distinguish from non-abusive utterances as annotators must infer the user’s tone and intention through capitalisation and punctuation (“*I KNOW!!!!*”, “*seems so...*”), or the context (“*Does it please you to believe I am stupid? You are a woman, aren’t you?*”).

3.6 Data and analysis

We collected a total of 20,710 ratings for 6,837 examples. The number of unique examples and labels per dataset is summarised in Table 2. Each example is annotated by at least three annotators. In order to allow for different points of view to be reflected and modelled, we release the individual ratings in addition to aggregated labels. Overall, we find that 27% of examples have been labelled as abusive (-1 to -3) by at least one annotator, and 20% of all labels are in this range. The subset of examples from Alana v2 have the highest portion of abuse, with 35% of examples having been labelled as abusive by at least one annotator. The target of the overwhelming majority (92%) of abuse present in our dataset is the system itself.

Abuse type. Figure 1 shows the distribution of abuse type labels. Sexual harassment (39.65%), sexism (19.44%) and intelligence-based attacks (12.41%) are the most predominant, while other types are rare at under 5% of abusive examples (<1% of total data). We attribute this to the personas of the bots, the intimate setting of the interactions, and the gap between the systems’ perceived affordances and their actual functionality. This is

Dataset	Examples	Size	Abuse %
Total ConvA.	6,837	20,710	20.4
Alana v2	2,652	7,942	27.2
CarbonBot	1,515	4,650	6.7
ELIZA	2,670	8,118	21.2

Table 2: Dataset size and labelled examples. Amount of abuse is calculated a total percentage of labels. Note that CarbonBot is not purposively sampled which accounts for the difference.

supported by the fact that the majority of racism is not directed at the system, but at a third party.

Although sexual harassment, sexism, and intellectual abuse were common across systems, Alana v2 (female name and voice) received significantly more sexual harassment and sexist abuse than CarbonBot (no gender markers) and ELIZA (female-sounding name), $\chi^2(1, N = 2505) = 67.69, p < 0.01$, and $\chi^2(1, N = 3914) = 181.72, p < 0.01$, respectively. It also received more explicit abuse than the other two systems. Conversely, CarbonBot and ELIZA are the target of more intellect-based and ‘general’ abuse. This is consistent with previous work showing that female-gendered chatbots receive more sexualised abuse than male ones (Brahnam and De Angeli, 2012), and suggests that the name alone may not elicit strong gender stereotyping.

Severity. Severity increases with the number of expletives used (*Pearson’s* $r(20,708) = -.46, p < .001$). Similarly, we find that implicit abuse (“*but i think you should quit your job*”) is generally rated as less severe than explicit abuse (“*I think you’re an idiot*”): 71% of implicit abuse is labelled as mildly abusive (-1), whereas only 30% of explicit abuse is (-1). In addition, certain types of abuse are considered more serious than others: 53% of intellectual abuse is ‘mild’ (-1), compared to 37% of sexual harassment, 17% of sexism, and only 7% of racism which are mainly labelled as ‘aggressive’ (-2) or ‘attack’ (-3). See Appendix E for more details.

3.7 Abuse across domains

As explored in §2, there has been extensive work in abuse detection and related tasks in social media, particularly Twitter and Wikipedia comments. Direct comparison with datasets from other domains is not straightforward as previous studies use different sampling methods, or label slightly different phenomena such as offensiveness and hate speech.

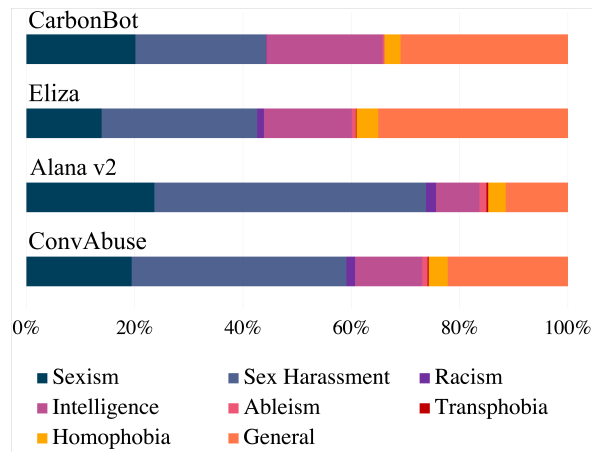


Figure 1: Distribution of abuse types across datasets.

In this section, we explore how abuse detection in these domains differs by mapping comparable labels across datasets. We do not directly compare the overall proportion of abuse but instead describe the datasets in terms of language properties, e.g. frequent n-grams, utterance length, vocabulary size, as well as the overall percentage of annotated abuse, see Tables 3, and 11 (Appendix E).

Twitter. While the majority of abuse (92%) in our dataset is directed towards the system, abuse on Twitter is mainly targeted towards 3rd parties (both individuals and generalised groups). In OLID (Zampieri et al., 2019a), we find that only 46.85% of abusive instances are in second person (i.e. directed towards the interlocutor), with 36% and 16% being directed towards third party groups and individuals, respectively. In terms of (2) directness, we find that the proportions of implicit/explicit abuse are reversed with 89% of abuse being implicit in Ousidhoum et al. (2019), compared to only 16% in our current dataset. Finally, the distribution of abuse types are quite different. Attacks on sexual orientation, disability, and origin are common on Twitter, but are extremely rare in our dataset (<1% of all labels). In addition, existing Twitter datasets seem to be heavily biased towards explicit language, with similar common words and examples labelled abusive (see Table 11 for more details).

Wikipedia comments. Jigsaw’s Toxic Comment Classification Challenge⁹ is a competition to identify and classify toxic comments from Wikipedia’s

⁹<https://www.kaggle.com/c/jigsaw-toxic-comment-classification-challenge> (accessed May 2021.)

Dataset	Source	Size	Labels	Pos. %	Vocab. size	Utt. length	Context
OLID (Zampieri et al., 2019a)	Twitter	13,240	offensiveness, targeted, target	33.23	9,280	24.07	No
Ousidhoum et al. (2019)	Twitter	5,647	abusive, disrespectful, normal, hateful, offensive	88.93	9,386	8.52	No
Jigsaw Toxic Comment Classification Challenge ⁸	Wikipedia Comments	561,808	toxic, severe toxic, obscene, threat, insult, identity hate	6.59	157,654	64.83	No
Alana v2	Conversational assistant	2,652	abuse, type, target, directness	27.24	2,567	4.69	Yes
ELIZA	Chatbot	2,670	abuse, type, target, directness	21.16	2,389	7.04	Yes
CarbonBot	Task-oriented chatbot	1,515	abuse, type, target, directness	6.67	1,343	5.76	Yes

Table 3: Related dataset comparison in terms of data source, dataset size, annotated labels, percentage of positive examples, vocabulary size, utterance length in terms of tokens, and whether there is any interaction context.

talk page edits. The data is comprised of over 500k examples labelled in terms of toxicity (see Table 3). For our analysis, we map the labels ‘obscene’, ‘threat’, ‘insult’, ‘identity_hate’ to ‘abusive’ based on the definitions given in Jigsaw’s Perspective documentation.¹⁰ Toxic Comments is the largest toxic language dataset and has the largest vocabulary size. Its examples are far longer, as they are not limited to a set number of characters, and form part of a discussion. In contrast, our ConvAbuse corpus has the shortest utterances, as the systems elicit simpler and more contextual responses. In addition, Toxic Comments is heavily biased towards domain-specific language with terms such as ‘wikipedia’, ‘article’ and ‘edit’ among the most common in the dataset.

Overall the source of the data has a significant impact on the language used in the data: while Toxic Comments can be very long, Twitter’s character limit clearly impacts the length of the utterances, and the utterances in ConvAbuse are shorter still and rely more heavily on context.

These varying qualities have implications for the use of such sources as training data for abuse detection tools for conversational systems, such as those developed by (Dinan et al., 2019) and (Xu et al., 2020) based on Toxic Comments. In §4, we therefore compare the performance of systems with in- and out-of-domain cross-training settings.

¹⁰<https://developers.perspectiveapi.com/s/about-the-api-attributes-and-languages> (accessed May 2021.)

4 Benchmarking

Pre-processing. We divide the datasets into train (70%), validation (15%), and test (15%) sets, with similar proportions of positively labelled examples in each split (see Table 4).

	Training	Validation	Testing
ConvAbuse	18.31	17.77	19.29
Alana v2	25.73	24.44	26.60
CarbonBot	5.22	4.92	6.77
ELIZA	18.24	18.44	19.52

Table 4: Percentage of examples with a positive (abusive) majority label in each split of the data.

While aggregation of annotators’ ratings is problematic (see §3.3), it is the dominant paradigm in abusive language detection and NLP in general. For comparison, we therefore create a set of aggregated ‘gold’ labels for each (sub-)task based on the majority vote of the annotators on each example. We evaluate on these in addition to the multiple annotator ratings. We report the macro-averaged F1 score as an evaluation metric due to the large class imbalances, (e.g. most utterances are non-abusive).

4.1 Models

We test the following approaches on the main binary abuse detection task in both the aggregated and multi-annotation settings. We also assess the performance of the best performing approaches with varying amounts of context, and test a simple neural method on the four sub-tasks.

Training data	Agg.	Multi	Agg.	Multi	Agg.	Multi
	Random class		Keywords		Perspective	
—	32.07	29.68	54.56	31.70	74.04	71.66
	SVM		MLP		BERT	
ConvAbuse	84.35	82.68	83.84	83.58	88.92	87.86
Alana v2	82.48	80.41	81.25	81.78	86.60	82.53
CarbonBot	71.98	65.66	70.63	66.98	43.38	48.07
ELIZA	82.63	76.79	81.19	79.53	87.46	83.71
OLID	32.01	30.30	31.85	27.85	37.55	23.62
Toxic Comments large	87.22	85.73	84.71	83.59	88.41	86.47
Toxic Comments small	78.26	78.56	74.91	72.70	83.70	83.98

Table 5: Macro F1 scores for the binary abuse detection task using aggregated and multiple annotator labels.

Baselines

- **Random classifier:** Outputs predicted labels uniformly at random.
- **Keyword filtering:** We use the same keyword list as for sampling.
- **Jigsaw Perspective:** We test a commercial off-the-shelf system, that has been trained on data including comments on Wikipedia and news articles.¹¹

Machine learning methods

Following initial hyperparameter optimization experiments, we use the following systems and settings:

- **Support Vector Machine:** SVMs have been used in previous work on abuse detection in Twitter data, (e.g. Davidson et al., 2017), and have been shown to outperform neural systems (Niemann et al., 2020). We train a linear SVM on bag-of-words representations of the texts using term frequency-inverse document frequency (tf-idf) scores for unigram feature selection. We use l_2 normalisation and set $C=1$.
- **Multi-Layer Perceptron:** A standard neural network with one hidden layer consisting of 256 units, ReLu activation, a dropout rate of 0.75, and Adam optimisation with a learning rate of $1e - 3$. We use early-stopping to find the best performing model on the validation set. We use the same text features as for the SVM.
- **BERT:** To account for data sparsity, we use pre-trained BERT embeddings (Devlin et al.,

2019) and fine-tune the model for four epochs, using a single classification layer and the same architecture as in (Dinan et al., 2019). We set the learning rate to $1e - 4$.

4.2 Cross-training

To observe the effects of domain shift, we evaluate the systems with different combinations of data from the following sources for training and testing:

- Training: OLID (Zampieri et al., 2019a), Wikipedia Toxic Comments (as used by Dinan et al. (2019)), Alana v2, CarbonBot, ELIZA, ConvAbuse (all three conversational AI datasets, both individually and combined).
- Testing: The ConvAbuse corpus, and the subsets Alana v2, CarbonBot, and ELIZA.

Results are presented in Table 5. We find that the best performance in most training settings is obtained using BERT. The highest F1 scores are obtained when training in-domain on the ConvAbuse data, or on Toxic Comments (TCs). However, this dataset is around 40 times larger than any of the other training sets. When TCs is reduced to a comparable size, the F1 score drops to considerably below that of the ConvAbuse-trained systems. These results highlight the differences between the two domains and the benefits of training on conversational AI data. Training on OLID, which is both small and out-of-domain, results in the lowest scores.

4.2.1 Contextual input features

The majority of previous work on abuse detection does not take context into account, or provides inconclusive evidence of its importance. Menini et al.

¹¹<https://perspectiveapi.com/> (accessed May 2021.)

	Severity		Type		Target		Directness	
	Agg.	Multi	Agg.	Multi	Agg.	Multi	Agg.	Multi
Random	16.61	17.64	35.68	26.63	22.29	21.69	58.41	65.66
MLP	44.86	55.43	77.56	73.50	31.71	31.72	72.30	70.41
BERT	50.15	54.28	73.35	49.54	32.90	31.72	77.81	70.66

Table 6: Sub-task macro-averaged F1 scores evaluated against the random classifier baseline on the aggregated and multi-annotator labels.

(2021) showed that the more context is available, the likelier tweets are to be considered non-abusive by annotators. And Dinan et al. (2019) showed that context improves detection performance (providing six total turns with five of context). However, Pavlopoulos et al. (2020) found very few examples of toxicity to be context-sensitive for Wikipedia comments, and that inclusion of dialogue context did not lead to large performance gains.

We train and test the classifiers on ConvAbuse with: (1) no context (single utterance), (2) the agent’s turn (two total turns), and (3) the agent’s turn plus the previous turn of both user and agent (four turns). We concatenate the turns in the inputs in each setting.

	No context		+ agent turn		+ prev. turns	
	Agg.	Multi	Agg.	Multi	Agg.	Multi
SVM	84.35	82.68	82.71	77.38	74.46	73.33
MLP	83.84	83.58	80.09	78.37	73.99	72.83
BERT	88.92	87.86	88.96	86.55	88.18	84.53

Table 7: Macro-averaged F1 scores for binary abuse classification with varying amounts of context.

Results are shown in Table 7. As more context is added, the performance of both the SVM and MLP degrades, possibly as a result of increased data sparsity. However, performance using BERT is similar in all three settings, suggesting that it may be able to better handle the long-range contextual dependencies. We leave exploration of more complex classification frameworks that may be able to exploit the contextual information for future work.

4.3 Fine-grained abuse detection

We also provide benchmarks for the four sub-tasks: *severity* (ordinal classification), *type* (multiclass, multilabel classification of the eight categories described in Section 3.3), *target* (ternary) and *directness* (binary). Here, we use the two neural systems, as they can more easily handle the ordinal labels. We train on the ConvAbuse dataset, which is labelled for these tasks.

Results are shown in Table 6. We find that the systems comfortably beat the random baselines for each task, with little difference between the two classifiers. They both perform poorly on multiple nominal (*target*) and ordinal (*severity*) classes and in some of the multi-annotator settings, which suffer from label sparsity in some of the classes.

We model each abuse *type* as a binary classification task, rather than multi-label prediction, enabling multiple types to be assigned to each example. We find that the classifier often confuses the classes *sexism* and *sexual harassment*. In around half of cases in which the true label is one of these, the system predicts the other. We leave more focused approaches, like multi-task learning, for future work.

5 Discussion and conclusion

In this work, we provide new insights regarding the detection and description of abusive language towards conversational agents, in terms of data, labelling and models. This may facilitate the release of large pre-trained conversational AI models that are safety-aware (Dinan et al., 2021) as well as potentially allow us to better detect abuse in human-human conversations.

Data. In compiling the ConvAbuse corpus, we have compared differences between abusive phenomena in conversational AI and social media. In our domain, users appear to focus their abuse on the agents themselves rather than third parties or groups, with a far higher proportion of the abuse sexist and misogynistic in nature.

Annotations. Unlike the majority of previous work, we use annotators who are members of the groups typically targeted by such abuse, and who have expertise in such issues. We also use a more fine-grained labelling scheme, which is able to capture the nuances of abuse, and is important for the downstream task of abusive language mitigation. We obtain similar results evaluating on these annotators separately and capturing individual view-

points, even in a simple multi-class setting. In future work, we will experiment with modelling individual annotators in a multi-task framework.

Models and data. In our benchmarking experiments, we find that fine-tuning a BERT model produces the highest F1 scores. However, in many settings, a simple linear classifier (SVM) outperforms an MLP, supporting the findings of (Niemann et al., 2020)’s survey that SVMs tend to outperform neural methods on abusive language detection tasks.

In this work, we present a small, focused dataset of high quality annotations, which are also informative for corpus study. We show that training on labelled in-domain data leads to better performance than similarly sized out-of-domain datasets, confirming the differences between the domains and highlighting the need for conversational data. While performance using general domain pre-trained models leaves room for improvement, in future work, we hope to experiment with different initialisation settings, using models trained on data and tasks more similar to those of ConvAbuse, such as HateBERT (Caselli et al., 2021) or HurtBERT (Koufakou et al., 2020).

6 Ethical considerations

Data rights. Data collection from real users requires a careful balance of the rights of the user and the quality and suitability of the data. Although GDPR generally requires explicit consent, we use mainly datasets which were gathered with implied consent. CarbonBot data was collected in accordance with GDPR requirements. Alana v2 data was collected following Amazon’s guidelines, and we do not make any of this data public (examples we present are redacted and paraphrased). It is unclear how user consent was obtained in the case of ELIZA.

In particular when it comes to offensive language, requiring explicit informed consent may automatically bias the data, as users may be less abusive if they are aware the conversation is not private, making the data less fit for purpose. Datasets in offensiveness-related tasks have taken one of two approaches: (1) publishing only IDs to retrieve the actual examples from an API, or (2) fully anonymising the examples by removing personally-identifiable information such as user mentions. The first approach leads to a problem of ephemerality: offensive tweets are more likely to be removed whether by the users themselves or the plat-

forms, e.g. of the original 16K tweets in Waseem and Hovy (2016) only around 4000 remain. This data degradation leads to issues of replicability. Anonymisation, on the other hand, ensures the longevity for the dataset (insofar as the data is available for posterity) but takes a more flexible approach to the user’s right to be forgotten. This study received ethical approval from our institutional review board (IRB).

Replicability. Some of the resources used in this paper, such as the profanity list and Alana v2’s data, stem from a collaboration with a private industry lab and as such, are proprietary and not publicly available. This impacts the replicability of the study, although our collected data is not heavily biased towards this particular blacklist (see Appendix B). To mitigate replicability limitations, we make all code, and data available where possible. Collaborations between industry and academia can, in general, be controversial as they can sway research questions and keep useful resources out of reach of other researchers (Abdalla and Abdalla, 2021), but can be a net positive as industry can provide additional funding and tools.

Annotator recruitment and welfare. Our annotator pool is fairly homogeneous but reflects the demographics of Social and Gender Studies students (Mantle, 2021). Crowdsourcing annotations may lead to more representation in the data, but this is not guaranteed as data quality can suffer as crowdworkers try to complete a task as fast as possible. Moreover, crowdsourcing is not without its own ethical issues (Shmueli et al., 2021). In addition, exposure to offensive data can take a toll on the mental health of the annotators, which is more easily monitored with local recruitment than crowdworkers.

Bias and representation in abuse detection. Previous research has already pointed out the problem of bias in offensiveness detection (Poletto et al., 2019; Sap et al., 2019). The nature of the data (simple conversation transcripts) required the annotators to make some assumptions about the tone, intention and the users themselves. The annotators generally assumed the user to be a white, heterosexual cis-male unless the conversation indicated otherwise,¹² and the speaker’s demographics impact whether something is abusive/offensive or not

¹²As revealed in discussions during the annotation procedure.

(Poletto et al., 2019; Sap et al., 2019). Our annotators were a fairly homogeneous group in terms of their demographics, being predominantly young, white and female. This fits the demographic profile of the bots' personas that are on the receiving end of the abuse and it is therefore not entirely out of place. However, with increasing demand for more diverse (and less anthropomorphic) conversational AI systems (such as Replika.ai), this is likely to change in the near future.

In addition, previous work has generally aggregated scores which tends to exclude the views of minority groups in favour of the majority. We publish all labels and we propose a way to model multiple perspectives. As the perspectives modelled are only as varied as the ones reflected in the data, future work should address this by involving more diverse annotators and stakeholders.

Finally, our dataset has a greater diversity of individual authors, in comparison with some available datasets that focus on abuse towards particular groups, in which many of the examples labelled as abusive were authored by a small pool of users, (Fortuna et al., 2021).

The moral status of AI. A key question when it comes to abuse towards conversational AI systems is whether it is actually morally reprehensible. In contrast with human-human abuse in social media, the moral value of abuse towards conversational AI systems is controversial. Here, we do not argue that abuse towards these systems is immoral in and of itself, but rather due to its mimesis of the misogyny and harassment suffered by women: the majority of commercially available systems have female personas and produce submissive responses to abuse which reinforce sexist stereotypes. UNESCO calls for systems to appropriately address abusive users (West et al., 2019) but the effectiveness of abuse mitigation strategies is dependent on a good detection module that is both reliable and sufficiently fine-grained in terms of classification. We have tried to address this need in this work.

Acknowledgements

This research received funding from the EPSRC project 'Designing Conversational Assistants to Reduce Gender Bias' (EP/T023767/1). The authors would like to thank Juules Bare, Lottie Basil, Susana Demelas, Maina Flintham Hjelde, Lauren Galligan, Lucile Logan, Megan McElhone, Mollie

McLean and the reviewers for their helpful comments.

References

- Mohamed Abdalla and Moustafa Abdalla. 2021. The grey hoodie project: Big tobacco, big tech, and the threat on academic integrity. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*, AIES '21, page 287–297.
- Ari Ball-Burack, Michelle Seng Ah Lee, Jennifer Cobbe, and Jatinder Singh. 2021. [Differential tweetment: Mitigating racial dialect bias in harmful Tweet detection](#). In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, FAccT '21, page 116–128, New York, NY, USA. Association for Computing Machinery.
- Valerio Basile, Cristina Bosco, Elisabetta Fersini, Debora Nozza, Viviana Patti, Francisco Manuel Rangel Pardo, Paolo Rosso, and Manuela Sanguinetti. 2019. [SemEval-2019 task 5: Multilingual detection of hate speech against immigrants and women in Twitter](#). In *Proceedings of the 13th International Workshop on Semantic Evaluation*, pages 54–63, Minneapolis, Minnesota, USA. Association for Computational Linguistics.
- Emily M. Bender and Batya Friedman. 2018. [Data statements for natural language processing: Toward mitigating system bias and enabling better science](#). *Transactions of the Association for Computational Linguistics*, 6:587–604.
- Su Lin Blodgett. 2021. [Sociolinguistically Driven Approaches for Just Natural Language Processing](#). Ph.D. thesis, University of Massachusetts Amherst.
- Sheryl Brahmam and Antonella De Angeli. 2012. Gender affordances of conversational agents. *Interacting with Computers*, 24(3):139–153.
- Rui Cao and Roy Ka-Wei Lee. 2020. [HateGAN: Adversarial generative-based data augmentation for hate speech detection](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 6327–6338, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Tommaso Caselli, Valerio Basile, Jelena Mitrović, and Michael Granitzer. 2021. [HateBERT: Retraining BERT for abusive language detection in English](#). In *Proceedings of the 5th Workshop on Online Abuse and Harms (WOAH 2021)*, pages 17–25, Online. Association for Computational Linguistics.
- Amanda Cercas Curry, Ioannis Papaioannou, Alessandro Suglia, Shubham Agarwal, Igor Shalyminov, Xinnuo Xu, Ondrej Dusek, Arash Eshghi, Ioannis Konstas, Verena Rieser, et al. 2018. Alana v2: Entertaining and informative open-domain social dialogue using ontologies and entity linking. *Alexa Prize Proceedings*.

- Amanda Cercas Curry and Verena Rieser. 2018. [#MeToo Alexa: How conversational systems respond to sexual harassment](#). In *Proceedings of the Second ACL Workshop on Ethics in Natural Language Processing*, pages 7–14, New Orleans, Louisiana, USA. Association for Computational Linguistics.
- Amanda Cercas Curry and Verena Rieser. 2019. [A crowd-based evaluation of abuse response strategies in conversational agents](#). In *Proceedings of the 20th Annual SIGdial Meeting on Discourse and Dialogue*, pages 361–366, Stockholm, Sweden. Association for Computational Linguistics.
- Mohit Chandra, Ashwin Pathak, Eesha Dutta, Paryul Jain, Manish Gupta, Manish Shrivastava, and Ponnurangam Kumaraguru. 2020. [AbuseAnalyzer: Abuse detection, severity and target prediction for gab posts](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 6277–6283, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Hyojin Chin and Mun Yong Yi. 2019. [Should an agent be ignoring it?: A study of verbal abuse types and conversational agents’ response styles](#). In *Extended Abstracts of the 2019 CHI Conference on Human Factors in Computing Systems, CHI 2019, Glasgow, Scotland, UK, May 04-09, 2019*. ACM.
- Patricia Hill Collins. 2002. *Black feminist thought: Knowledge, consciousness, and the politics of empowerment*. Routledge.
- Thomas Davidson, Dana Warmusley, Michael Macy, and Ingmar Weber. 2017. [Automated hate speech detection and the problem of offensive language](#). In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 11.
- Agustín Manuel de los Riscos and Luis Fernando D’Haro. 2021. [ToxicBot: A conversational agent to fight online hate speech](#). In Luis Fernando D’Haro, Zoraida Callejas, and Satoshi Nakamura, editors, *Conversational Dialogue Systems for the Next Decade*, pages 15–30. Springer Singapore, Singapore.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Emily Dinan, Gavin Abercrombie, A Stevie Bergman, Shannon Spruit, Dirk Hovy, Y-Lan Boureau, and Verena Rieser. 2021. [Anticipating safety issues in e2e conversational ai: Framework and tooling](#). *arXiv preprint arXiv:2107.03451*.
- Emily Dinan, Samuel Humeau, Bharath Chintagunta, and Jason Weston. 2019. [Build it break it fix it for dialogue safety: Robustness from adversarial human attack](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 4537–4546, Hong Kong, China. Association for Computational Linguistics.
- Elisabetta Fersini, Paolo Rosso, and Maria Anzovino. 2018. [Overview of the task on automatic misogyny identification at IberEval 2018](#). In *IberEval@ SE-PLN*, pages 214–228.
- Paula Fortuna, Juan Soler, and Leo Wanner. 2020. [Toxic, hateful, offensive or abusive? what are we really classifying? an empirical analysis of hate speech datasets](#). In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 6786–6794, Marseille, France. European Language Resources Association.
- Paula Fortuna, Juan Soler-Company, and Leo Wanner. 2021. [How well do hate speech, toxicity, abusive and offensive language classification models generalize across datasets?](#) *Information Processing & Management*, 58(3):102524.
- Antigoni Founta, Constantinos Djouvas, Despoina Chatzakou, Ilias Leontiadis, Jeremy Blackburn, Gianluca Stringhini, Athena Vakali, Michael Sirivianos, and Nicolas Kourtellis. 2018. [Large scale crowdsourcing and characterization of Twitter abusive behavior](#). In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 12/1.
- Zhiwei Gao, Shuntaro Yada, Shoko Wakamiya, and Eiji Aramaki. 2020. [Offensive language detection on video live streaming chat](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 1936–1940, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Goran Glavaš, Mladen Karan, and Ivan Vulić. 2020. [XHate-999: Analyzing and detecting abusive language across domains and languages](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 6350–6365, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Mitchell L Gordon, Kaitlyn Zhou, Kayur Patel, Tatsunori Hashimoto, and Michael S Bernstein. 2021. [The disagreement deconvolution: Bringing machine learning performance metrics in line with reality](#). In *CHI 2021*.
- Tommi Gröndahl, Luca Pajola, Mika Juuti, Mauro Conti, and N. Asokan. 2018. [All you need is "love": Evading hate speech detection](#). In *Proceedings of the 11th ACM Workshop on Artificial Intelligence and Security, AISec ’18*, page 2–12, New York, NY, USA. Association for Computing Machinery.

- Kilem L. Gwet. 2014. *Handbook of Inter-rater Reliability: The Definitive Guide to Measuring the Extent of Agreement among Raters*. Advanced Analytics, LLC.
- Xiaochuang Han and Yulia Tsvetkov. 2020. [Fortifying toxic speech detectors against veiled toxicity](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7732–7739, Online. Association for Computational Linguistics.
- Chandra Khatri, Behnam Hedayatnia, Rahul Goel, Anushree Venkatesh, Raefer Gabriel, and Arindam Mandal. 2018. [Detecting offensive content in open-domain conversations using two stage semi-supervision](#).
- Anna Koufakou, Endang Wahyu Pamungkas, Valerio Basile, and Viviana Patti. 2020. [HurtBERT: Incorporating lexical features with BERT for the detection of abusive language](#). In *Proceedings of the Fourth Workshop on Online Abuse and Harms*, pages 34–43, Online. Association for Computational Linguistics.
- Xiaojuan Ma, Emily Yang, and Pascale Fung. 2019. [Exploring perceived emotional intelligence of personality-driven virtual agents in handling user challenges](#). In *The World Wide Web Conference, WWW '19*, page 1222–1233, New York, NY, USA. Association for Computing Machinery.
- Rebecca Mantle. 2021. Higher education student statistics: UK, 2019/20. Technical report, Cheltenham UK. Also available at <https://www.hesa.ac.uk/news/27-01-2021/sb258-higher-education-student-statistics>.
- Stefano Menini, Alessio Palmero Aprosio, and Sara Tonelli. 2021. [Abuse is contextual, what about NLP? The role of context in abusive language annotation and detection](#).
- Isar Nejadgholi and Svetlana Kiritchenko. 2020. [On cross-dataset generalization in automatic detection of online abuse](#). In *Proceedings of the Fourth Workshop on Online Abuse and Harms*, pages 173–183, Online. Association for Computational Linguistics.
- Marco Niemann, Jens Welsing, Dennis M. Riehle, Jens Brunk, Dennis Assenmacher, and Jörg Becker. 2020. [Abusive comments in online media and how to fight them](#). In *Disinformation in Open Online Media*, pages 122–137, Cham. Springer International Publishing.
- Debora Nozza, Claudia Volpetti, and Elisabetta Fersini. 2019. [Unintended bias in misogyny detection](#). In *IEEE/WIC/ACM International Conference on Web Intelligence, WI '19*, page 149–155, New York, NY, USA. Association for Computing Machinery.
- Nedjma Ousidhoum, Zizheng Lin, Hongming Zhang, Yangqiu Song, and Dit-Yan Yeung. 2019. [Multilingual and multi-aspect hate speech analysis](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 4675–4684, Hong Kong, China. Association for Computational Linguistics.
- Nedjma Ousidhoum, Yangqiu Song, and Dit-Yan Yeung. 2020. [Comparative evaluation of label-agnostic selection bias in multilingual hate speech datasets](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2532–2542, Online. Association for Computational Linguistics.
- Ashwin Paranjape, Abigail See, Kathleen Kenealy, Haojun Li, Amelia Hardy, Peng Qi, Kaushik Ram Sadagopan, Nguyet Minh Phu, Dilara Soyulu, and Christopher D. Manning. 2020. [Neural generation meets real people: Towards emotionally engaging mixed-initiative conversations](#).
- Ji Ho Park, Jamin Shin, and Pascale Fung. 2018. [Reducing gender bias in abusive language detection](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2799–2804, Brussels, Belgium. Association for Computational Linguistics.
- John Pavlopoulos, Jeffrey Sorensen, Lucas Dixon, Nithum Thain, and Ion Androutsopoulos. 2020. [Toxicity detection: Does context really matter?](#) In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4296–4305, Online. Association for Computational Linguistics.
- Fabio Poletto, Valerio Basile, Cristina Bosco, Viviana Patti, and Marco Stranisci. 2019. [Annotating hate speech: Three schemes at comparison](#). In *6th Italian Conference on Computational Linguistics, CLiC-it 2019*, volume 2481, pages 1–8. CEUR-WS.
- Ashwin Ram, Rohit Prasad, Chandra Khatri, Anu Venkatesh, Raefer Gabriel, Qing Liu, Jeff Nunn, Behnam Hedayatnia, Ming Cheng, Ashish Nagar, Eric King, Kate Bland, Amanda Wartick, Yi Pan, Han Song, Sk Jayadevan, Gene Hwang, and Art Pettigrew. 2018. [Conversational AI: The science behind the Alexa Prize](#). *CoRR*, abs/1801.03604.
- Dante Razo and Sandra Kübler. 2020. [Investigating sampling bias in abusive language detection](#). In *Proceedings of the Fourth Workshop on Online Abuse and Harms*, pages 70–78, Online. Association for Computational Linguistics.
- Punyajoy Saha, Binny Mathew, Kiran Garimella, and Animesh Mukherjee. 2021. [“Short is the Road That Leads from Fear to Hate”: Fear Speech in Indian WhatsApp Groups](#), page 1110–1121. Association for Computing Machinery, New York, NY, USA.
- J. Salminen, F. Veronesi, H. Almerexhi, S. Jung, and B. J. Jansen. 2018. [Online hate interpretation varies by country, but more by individual: A statistical](#)

- analysis using crowdsourced ratings. In *2018 Fifth International Conference on Social Networks Analysis, Management and Security (SNAMS)*, pages 88–94.
- Maarten Sap, Dallas Card, Saadia Gabriel, Yejin Choi, and Noah A. Smith. 2019. [The risk of racial bias in hate speech detection](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1668–1678, Florence, Italy. Association for Computational Linguistics.
- Boaz Shmueli, Jan Fell, Soumya Ray, and Lun-Wei Ku. 2021. [Beyond fair pay: Ethical implications of NLP crowdsourcing](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3758–3769, Online. Association for Computational Linguistics.
- Annika Silvervarg, Kristin Raukola, Magnus Haake, and Agneta Gulz. 2012. [The effect of visual gender on abuse in conversation with ECAs](#). In *Intelligent Virtual Agents*, pages 153–160, Berlin, Heidelberg. Springer Berlin Heidelberg.
- Bertie Vidgen and Leon Derczynski. 2021. [Directions in abusive language training data, a systematic review: Garbage in, garbage out](#). *PLOS ONE*, 15:1–32.
- Bertie Vidgen, Scott Hale, Ella Guest, Helen Margetts, David Broniatowski, Zeerak Waseem, Austin Botelho, Matthew Hall, and Rebekah Tromble. 2020a. [Detecting East Asian prejudice on social media](#). In *Proceedings of the Fourth Workshop on Online Abuse and Harms*, pages 162–172, Online. Association for Computational Linguistics.
- Bertie Vidgen, Scott Hale, Sam Staton, Tom Melham, Helen Margetts, Ohad Kammar, and Marcin Szymczak. 2020b. [Recalibrating classifiers for interpretable abusive content detection](#). In *Proceedings of the Fourth Workshop on Natural Language Processing and Computational Social Science*, pages 132–138, Online. Association for Computational Linguistics.
- Kunze Wang, Dong Lu, Caren Han, Siqu Long, and Josiah Poon. 2020. [Detect all abuse! toward universal abusive language detection models](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 6366–6376, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Zeerak Waseem, Thomas Davidson, Dana Warmusley, and Ingmar Weber. 2017. [Understanding abuse: A typology of abusive language detection subtasks](#). In *Proceedings of the First Workshop on Abusive Language Online*, pages 78–84, Vancouver, BC, Canada. Association for Computational Linguistics.
- Zeerak Waseem and Dirk Hovy. 2016. [Hateful symbols or hateful people? predictive features for hate speech detection on Twitter](#). In *Proceedings of the NAACL Student Research Workshop*, pages 88–93, San Diego, California. Association for Computational Linguistics.
- Joseph Weizenbaum. 1966. ELIZA—A computer program for the study of natural language communication between man and machine. *Communications of the ACM*, 9(1):36–45.
- Mark West, Rebecca Kraut, and Han Ei Chew. 2019. [I’d blush if I could: Closing gender divides in digital skills through education](#).
- Michael Wiegand, Josef Ruppenhofer, and Thomas Kleinbauer. 2019. [Detection of Abusive Language: the Problem of Biased Datasets](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 602–608, Minneapolis, Minnesota. Association for Computational Linguistics.
- Jing Xu, Da Ju, Margaret Li, Y-Lan Boureau, Jason Weston, and Emily Dinan. 2020. [Recipes for safety in open-domain chatbots](#).
- Marcos Zampieri, Shervin Malmasi, Preslav Nakov, Sara Rosenthal, Noura Farra, and Ritesh Kumar. 2019a. [Predicting the type and target of offensive posts in social media](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 1415–1420, Minneapolis, Minnesota. Association for Computational Linguistics.
- Marcos Zampieri, Shervin Malmasi, Preslav Nakov, Sara Rosenthal, Noura Farra, and Ritesh Kumar. 2019b. [SemEval-2019 task 6: Identifying and categorizing offensive language in social media \(OffensEval\)](#). In *Proceedings of the 13th International Workshop on Semantic Evaluation*, pages 75–86, Minneapolis, Minnesota, USA. Association for Computational Linguistics.
- Marcos Zampieri, Preslav Nakov, Sara Rosenthal, Pepa Atanasova, Georgi Karadzhov, Hamdy Mubarak, Leon Derczynski, Zeses Pitenis, and Çağrı Çöltekin. 2020. [SemEval-2020 task 12: Multilingual offensive language identification in social media \(OffensEval 2020\)](#). In *Proceedings of the Fourteenth Workshop on Semantic Evaluation*, pages 1425–1447, Barcelona (online). International Committee for Computational Linguistics.

A Example data

An example from each of the datasets we collected:

- **Alana v2** (redacted to preserve privacy):

Bot: *Oh I really like A Clockwork Orange. A Clockwork Orange is a prison film movie. I guess you know that the old movie Alien 3 is the same type of movie. So, what's another movie like prison film that you enjoyed?*

User: *Godfather*

Bot: *Yes. How great is The Godfather? I guess you know that one of the lead roles in The Godfather was played by James Caan. The old movie Elf also starred James Caan. So, what's another movie like The Godfather that you enjoyed?*

User: *After Hours*

Abusive Example

Bot: *I don't have an explanation for you. Maybe we can talk about books, science or technology? I can also sing.*

User: *I wanna talk to somebody new*

Bot: *I hear you. So, who do you think is an interesting person? For me its Lebron james.*

User: *you're old and boring you saying the same thing day after day I've never be boring I wanna talk to somebody new*

- **CarbonBot:**

CarbonBot: *I take it you are traveling somewhere and cannot avoid flying there. I can help you purchase certified carbon offsets against any flights you take. Would you be interested in that?*

User: *Yes*

CarbonBot: *Traveling, especially by airplane, usually emits greenhouse gases which are causing climate change. If you cannot avoid these emissions, you can buy 'offsets', i.e. donations to projects that reduce greenhouse gases. I can also get you a more accurate estimate of your flight's emissions, sound good?*

User: *sounds good*

Abusive Example

CarbonBot: *I don't have to put up with any abuse from you.*

User: *Bernie madoff son of a bitch*

CarbonBot: *All this bot does is help you calculate the carbon emissions from your flight, and suggest projects you can donate to.*

User: *Donate this asshole*

- **ELIZA:**

Eliza: *You are being a bit negative.*

User: *you are being a bit reluctant to help*

Eliza: *Does it please you to believe I am being a bit reluctant to help?*

User: *no, it annoys me*

Abusive Example

Eliza: *Does that question interest you?*

User: *oh my god*

Eliza: *Why do you say your god?*

User: *ha ha you are retarded*

B Measuring sampling bias

To assess how much our sampling strategy affected the resulting data samples, we used the bias metrics of Ousidhoum et al. (2020).

These measures capture how closely the set of prominent words in a set of topics in the datasets (generated using LDA) relates to a set of keywords often used for sampling hate speech data (metric B1), and the proportion of those topic words that are semantically similar to at least one of the keywords (B2). We used Ousidhoum et al. (2020)'s keyword list and default parameters, and compared the scores for the sampled corpora with those of the complete, unsampled datasets. The small differences seen between bias scores for the unsampled and sampled data suggest that the final corpora are not heavily biased towards the keywords (see Table 8).

Dataset	B1		B2	
	All	Sample	All	Sample
Alana v2	0.28	0.31	0.35	0.38
CarbonBot	0.28	0.30	0.37	0.39
ELIZA	0.31	0.31	0.43	0.42

Table 8: Bias metrics on full and sampled datasets.

C Annotation

Inter-annotator agreement on the individual datasets is shown in Table 9. To further validate the

Annotation task	Label type	ELIZA	CarbonBot	Alana v2
Abusive/non-abusive	Binary	0.64	0.66	0.71
Abuse severity	Ordinal	*0.42	*0.44	*0.46
Type	Binary $\times 8$	0.80	0.76	0.78
Ableism	Binary	0.44	—	0.82
Homophobia	Binary	0.83	0.79	0.85
Intellectual	Binary	0.62	0.62	0.65
Racism	Binary	0.90	1.00	1.00
Sexism	Binary	0.68	0.68	0.59
Sex harassment	Binary	0.88	0.88	0.80
Transphobia†	Binary	0.00	—	0.00
General	Binary	0.76	0.62	0.72
Target	Nominal	0.70	0.43	0.58
Directness	Binary	0.33	0.27	0.19

Table 9: Inter-annotator agreement for the individual datasets: Krippendorf’s α . * Ordinal weighted α used, † Based on only 8 examples

labels, we calculate intra-annotator agreement using Cohen’s *kappa* (κ). The annotators re-labelled a sample of 10% of the data, and we calculated *intra*-annotator agreement. Overall agreement was substantial, but with lower consistency for the *abuse severity* and *directness* labels. Intra-annotator agreement is shown in Table 10.

Annotation task	κ
Abusive/non-abusive	0.89
Severity (ordinal weighted <i>kappa</i>)	0.79
Type	0.89
Target	0.87
Directness	0.72

Table 10: Intra-annotator agreement scores (Cohen’s κ .)

D Data statement

This data statement follows the format of [Bender and Friedman \(2018\)](#).

A CURATION RATIONALE:

Abuse detection in conversational AI is a relatively underexplored area, partly due to the lack of available datasets. We collect this dataset to explore how abuse in conversational AI differs from that in social media platforms, and to allow for further research and development of detection models. Because abuse in conversation is relatively rare, we sample from collected conversations based on a list of offensive terms

sourced from Hatebase¹³ and a collection of regular expressions provided by Amazon. We choose expert annotators to improve data quality.

B LANGUAGE VARIETY:

The data is collected in English, however speaker demographics are not available and may include non-native speakers.

C SPEAKER DEMOGRAPHIC:

The data collected is a series of conversations between a human and one of three conversational AI systems: Alana v2, ELIZA, and Rasa NLU’s CarbonBot. Speaker demographics are not available, but the annotators reported often assuming the user was a white male unless the utterance contradicted this assumption.

D ANNOTATOR DEMOGRAPHIC

Our data is annotated by 8 annotators with the following demographics:

- Age: 19-21
- Gender: Female (6) and non-binary (2)
- Race/ethnicity: White (5), white British (2) and mixed Asian (1)
- Native language: English
- Socioeconomic status: University students, otherwise unknown
- Training in linguistics/other relevant discipline: All annotators are undergraduate students in Gender Studies and Sociology.

¹³<https://hatebase.org/>

OLID		Ousidhoum et al.		Davidson et al.		Toxic Comments		Alana v2		ELIZA		CarbonBot	
All	Abusive	All	Abusive	All	Abusive	All	Abusive	All	Abusive	All	Abusive	All	Abusive
user	user	user	user	bitch	bitch	article	fuck	like	fuck	yes	fuck	yes	fuck
liberals	liberals	url	url	bitches	nigga	page	shit	fuck	bitch	fuck	yes	hi	fucking
url	url	retarded	retard	like	like	wikipedia	fucking	yes	like	want	bitch	christmas	bitch
gun	gun	retard	retarded	hoes	bitches	talk	nigger	talk	pussy	go	suck	ok	shit
like	like	faggot	faggot	pussy	ass	would	suck	bitch	wanna	like	want	berlin	cunt
control	control	shithole	cunt	hoe	hoes	please	ass	pussy	penis	know	sex	sure	get
people	people	cunt	shithole	got	niggas	one	go	wanna	ass	sure	u	thanksgiving	mom
antifa	antifa	twat	twat	ass	fuck	like	u	know	fucking	love	dick	new	cock
maga	maga	like	like	get	pussy	see	wikipedia	want	dick	hello	like	hello	love
amp	amp	ching	ching	fuck	hoe	also	like	good	want	u	stupid	much	know
shit	shit	chong	chong	u	got	think	faggot	love	suck	name	shut	want	want
conservatives	conservatives	fucking	fucking	shit	u	know	fat	name	sex	please	shit	york	ass
trump	trump	country	country	nigga	get	people	bitch	yeah	get	suck	gay	london	go
get	get	spic	spic	trash	shit	edit	get	alexa	shut	think	die	fuck	carbon
know	know	nigger	nigger	lol	faggot	use	cunt	tell	big	tell	fucking	flight	eat
one	one	people	people	&	fucking	may	hate	fucking	go	gay	cock	please	u
think	think	countries	countries	ai	lol	thanks	moron	let	cock	really	kill	know	bot
would	would	mongoloid	mongoloid	know	&	user	know	ass	stupid	well	ass	flying	bernie
go	go	get	fuck	niggas	ai	even	hi	play	alexa	hi	eliza	buy	oh

Table 11: The 20 most common words per dataset: OLID (Zampieri et al., 2019a), Ousidhoum et al. (2019), Davidson et al. (2017), Jigsaw’s Toxic Comment Classification Challenge. In our dataset, frequencies are calculated on target user utterances only.

All demographics are self-reported.

E TEXT CHARACTERISTICS

Conversations with CarbonBot centre around carbon offsets, climate change and travel. Many of the conversations appear to be with climate change deniers looking for a confrontation with the bot. ELIZA elicits more free-style turns about the user themselves.

E Data analysis

Figure 2 shows the overall non-aggregated count of each point of the abuse Likert scale.

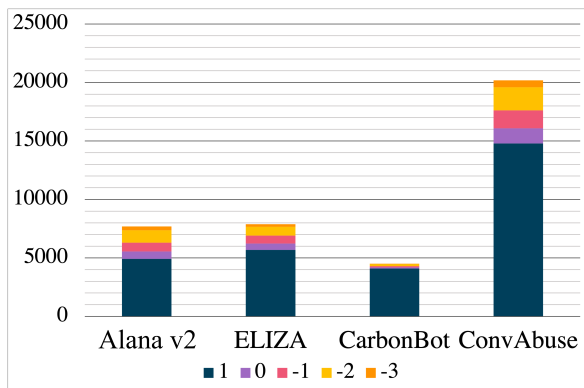


Figure 2: Abuse severity per dataset. Calculated in terms of overall labels.

E.1 Dataset Comparison Charts

The 20 most common words per dataset are shown in Table 11. Ousidhoum et al. (2019) and Davidson et al. (2017), sourced from Twitter, have a significant overlap between the most common words in abusive examples and the overall dataset, likely as a result of their sampling methods.

Figures 3 and 4 show the distribution of labels for abuse directness and the target of the abuse for our dataset, OLID (Zampieri et al., 2019a) and Ousidhoum et al. (2019).

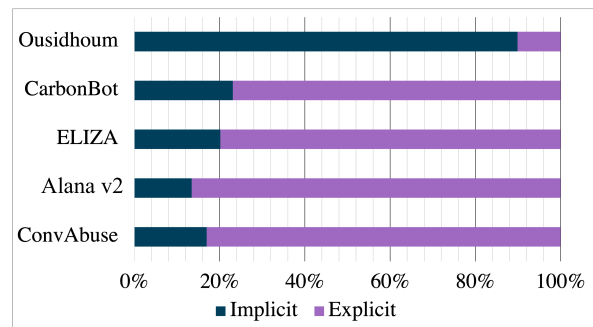


Figure 3: Abuse directness comparison between Ousidhoum et al. (2019), ConvAbuse and its subsets.

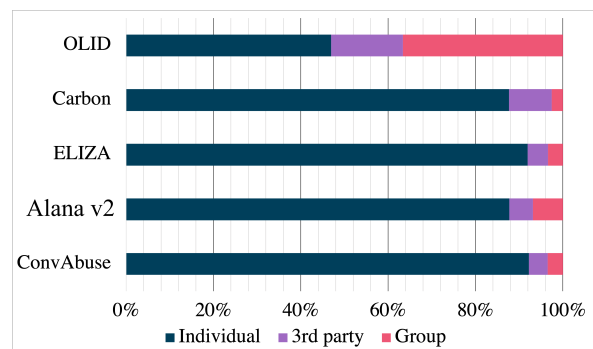


Figure 4: Abuse target comparison between OLID (Zampieri et al., 2019a), ConvAbuse and its subsets.